

Efficiency and Budget Balance in General Quasi-linear Domains ^{*}

Swaprava Nath¹ and Tuomas Sandholm¹

¹Carnegie Mellon University, Computer Science Department, {swapravn,sandholm}@cs.cmu.edu

Abstract

We study *efficiency* and *budget balance* for designing mechanisms in general quasi-linear domains. Green and Laffont (1979) proved that one cannot generically achieve both. We consider strategyproof budget-balanced mechanisms that are approximately efficient. For deterministic mechanisms, we show that a strategyproof and budget-balanced mechanism must have a *sink* agent whose valuation function is ignored in selecting an alternative, and she is compensated with the payments made by the other agents. We assume the valuations of the agents come from a bounded open interval. This result strengthens Green and Laffont’s impossibility result by showing that even in a restricted domain of valuations, there does not exist a mechanism that is strategyproof, budget balanced, and takes every agent’s valuation into consideration—a corollary of which is that it cannot be efficient. Using this result, we find a tight lower bound on the inefficiencies of strategyproof, budget-balanced mechanisms in this domain. The bound shows that the inefficiency asymptotically disappears when the number of agents is large—a result close in spirit to Green and Laffont (1979, Theorem 9.4). However, our results provide worst-case bounds and the best possible rate of convergence.

Next, we consider minimizing any convex combination of inefficiency and budget imbalance. We show that if the valuations are unrestricted, no deterministic mechanism can do asymptotically better than minimizing inefficiency alone.

Finally, we investigate randomized mechanisms and provide improved lower bounds on expected inefficiency. We give a tight lower bound for an interesting class of strategyproof, budget-balanced, randomized mechanisms. We also use an optimization-based approach—in the spirit of *automated mechanism design*—to provide a lower bound on the minimum achievable inefficiency of any randomized mechanism.

Experiments with real data from two applications show that the inefficiency for a simple randomized mechanism is 5–100 times smaller than the worst case. This relative difference increases with the number of agents.

1 Introduction

Consider a group of friends deciding which movie to watch together. Furthermore, the movie can be watched in someone’s home by renting it or at any of a number of movie theaters. Each of these choices incurs a cost. Since individual preferences are different and sometimes conflicting,

^{*}We thank Ioannis Caragiannis, Vincent Conitzer, Debasis Mishra, Hervé Moulin, Ariel Procaccia, and Arunava Sen for useful discussions. Nath is funded by the Fulbright-Nehru postdoctoral fellowship. Sandholm is funded by the National Science Foundation under grants 1320620 and 1546752.

the final choice may not make everybody maximally satisfied. This may cause some of the agents to misreport their preferences or drop out of the plan. To alleviate this problem, one can think of monetary transfers so friends who get their more-preferred choice pay more than friends that get their less-preferred choice. Desirable properties of such a choice and payment rule are that (1) the total side payments (transfers among the friends) should sum to zero, so there is no surplus or deficit, and (2) the choice is efficient, that is, the movie that is selected maximizes the sum of all the friends' valuations. Since the valuations are private information of the friends, an efficient decision requires the valuations to be revealed truthfully. This simple example problem is representative of many joint decision-making problems that often involve monetary transfers. Consider, for example, a group of firms sharing time on a jointly-owned supercomputer, city dwellers deciding on the location and choice of a public project (e.g., stadium, subway, or library), mobile service providers dividing spectrum among themselves, or a student body deciding which musician or art performer to invite to entertain at their annual function. These problems all call for efficient joint decision making and involve—or could involve depending on the application—monetary transfers.

This is a ubiquitous problem in practice and a classic problem in the academic literature. We study the standard model of this problem where the agents' utilities are *quasi-linear*: each agent's utility is her valuation for the selected alternative (e.g., the choice of movie) minus the money she has to pay. The classic goal is to select an *efficient* alternative, that is, the one that maximizes the sum of the agents' valuations (also known as *social welfare*).

In the setting where valuations are private information, a mechanism needs to be designed that incentivizes the agents to reveal their valuations truthfully (by the revelation principle, there is no loss in objective from restricting attention to such direct-revelation mechanisms). We will study the problem of designing strategyproof mechanisms, that is, mechanisms where each agent is best off revealing the truth regardless of what other agents reveal.

Even though there are mechanisms that select efficient alternatives in a truthful manner (e.g., the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973)), the transfers by the individuals do not sum to zero (in public goods settings, the VCG mechanism leads to too much money being collected from the agents). The execution of such a mechanism needs an external mediator who consumes the surplus (or may need to pay the deficit), to keep the mechanism truthful and efficient—a phenomenon known as ‘money burning’ in literature. In our movie selection example, this implies that we need a third party who will collect the additional money paid by the individuals, which is highly impractical in many settings. This has attracted significant criticism of the VCG mechanism (Rothkopf, 2007). Ideally, one would like to design strategyproof mechanisms that are efficient and *budget balanced*, that is, they do not have any surplus or deficit. Green and Laffont (1979) proved a seminal impossibility for this setting: in the general quasi-linear domain, strategyproof, efficient mechanisms cannot be budget balanced.

In this paper, we primarily focus on the problem of minimizing inefficiency subject to budget balance in the general setting of quasi-linear utilities. This is because, in the applications of interest to this paper (e.g., movie selection), budget balance is more critical than efficiency. However, we show that for a large set of agents, the per-agent inefficiency vanishes. We also show that for deterministic settings, optimizing the sum (or any convex combination) of efficiency and budget balance—which seems to be the most sensible objective—does not provide any asymptotic benefit over maximizing efficiency subject to budget balance. The main contributions of this paper are summarized in the following subsection.

1.1 Contributions of this paper

In this paper, we assume that the agents’ valuations are picked from a bounded open interval. In Section 3, we characterize the structure of truthful, budget balanced, *deterministic* mechanisms in this restricted domain, and show that any such mechanism must have a *sink* agent,¹ whose reported valuation function does not impact the choice of alternative and she gets the payments made by the other agents (Footnote 5). This result strengthens the Green and Laffont impossibility by showing that even in a restricted domain of bounded valuations, there does not exist a mechanism that is strategyproof, budget balanced, and takes every agent’s valuation into consideration—a corollary of which is that it cannot be efficient. With the help of this characterization, we find the optimal deterministic mechanism that minimizes the inefficiency. This provides a tight lower bound on the inefficiency of deterministic, strategyproof, budget-balanced mechanisms. By inefficiency of a mechanism in this paper, we mean the worst-case inefficiency over all valuation profiles. We provide a precise rate of decay ($\frac{1}{n}$) of the inefficiency with the increase in the number of agents (Theorem 2). This implies that the inefficiency vanishes for large number of agents.

To contrast this mechanism with the class of mechanisms that minimize budget imbalance subject to efficiency, in Section 4 we consider the joint objective of *efficiency-budget spillover*, which is a convex combination of inefficiency and budget imbalance. We prove that if the valuations are unrestricted, no deterministic, strategyproof mechanism can reduce this spillover at a rate faster than $\frac{1}{n}$ (Theorem 3). In other words, in the deterministic setting, minimizing the joint objective does not give any asymptotic advantage over the solution of minimizing inefficiency with the constraint that the mechanism is budget balanced.

We investigate the advantages of randomized mechanisms in Section 5. We first consider the class of *generalized sink* mechanisms. These mechanisms have, for every possible valuation profile, a probability distribution over the agents that determines each agent’s chance of becoming the sink. This class of mechanisms is budget balanced by design. We show examples where mechanisms from this class are not strategyproof (Algorithm 2), and then isolate an interesting subclass whose mechanisms are strategyproof, the *modified irrelevant sink mechanisms* (Algorithm 3). We show that no mechanism from this class can perform better than the deterministic mechanisms if the number of alternatives is greater than the number of agents (Theorem 4). Since inefficiency (weakly) increases with the number of alternatives (Theorem 5), we consider the extreme case of two alternatives and compare the performances of different mechanisms. We show that a naïve uniform random sink mechanism and the modified irrelevant sink mechanism (Algorithm 3) perform equally well (Theorems 6 and 7) and reduce the inefficiency by a constant factor of 2 from that of the deterministic mechanisms. However, the optimal, strategyproof, budget-balanced, randomized mechanism performs better than these mechanisms. Since the structure of strategyproof randomized mechanisms for general quasi-linear utilities is unknown,² we take a computational optimization-based approach to find the best mechanism for the special case of two agents. This approach is known in

¹Mechanisms using this idea have been presented with different names in the literature. The original paper by Green and Laffont (1979) refers to this kind of agents as a *sample* of the population. Later Gary-Bobo and Jaaidane (2000) formalized the randomized version of this mechanism which is known as *polling* mechanism. Faltings (2004) refers to this as an *excluded coalition* (when there are multiple such agents) and Moulin (2009) mentions this as *residual claimants*. However, we use the term ‘sink’ for brevity and convenience, and our paper considers a different setup and optimization objective.

²For randomized mechanisms, results involving special domains are known, e.g., facility location (Thang, 2010; Procaccia and Tennenholtz, 2009; Feldman and Wilf, 2011), auctions (Dobzinski et al., 2006), kidney exchange (Ashlagi et al., 2013), and most of these mechanisms aim for specific objectives.

the literature as *automated mechanism design* (Conitzer and Sandholm, 2002). For an overview, see Sandholm (2003). We prove that for a discrete valuation space with 3 levels, the optimal inefficiency is reduced by a factor of 7 (Theorem 9) from that of deterministic mechanisms. However, when the number of levels increases—thereby making the lower bound tighter to the actual open-interval problem—the improvement factor reduces to less than 5 (Figure 1). This is a significant improvement over the class of randomized sink mechanisms, which only improve over the best deterministic mechanism by a factor of 2.

We present experiments using real data from two applications. They show that in practice the inefficiency is significantly smaller than the worst case bounds (Section 6). We conclude the paper in Section 7 and present future research directions. Owing to the page limitation, we defer most of the proofs to the appendix.

1.2 Relationship to the literature

The Green-Laffont impossibility result motivated the research direction of designing efficient mechanisms that are minimally budget imbalanced. The approach is to redistribute the surplus money in a way that satisfies truthfulness of the mechanism (Bailey, 1997; Cavallo, 2006). The *worst case optimal* and *optimal in expectation* guarantees have been given for this class of mechanisms in restricted settings (Guo and Conitzer, 2008; Moulin, 2009; Guo and Conitzer, 2009). The performance of this class of *redistribution* mechanisms has been evaluated in interesting special domains such as allocating single or multiple (identical or heterogeneous) objects (Gujar and Narahari, 2011). Also, mechanisms have been developed and analyzed that are budget balanced (or no deficit) and minimize the inefficiency in special settings (Massó et al., 2015; Guo and Conitzer, 2014; Mishra and Sharma, 2016). Characterization of strategyproof budget-balanced mechanisms in the setting of cost-sharing is explored by Moulin and Shenker (2001) and its quantitative guarantees are presented by Roughgarden and Sundararajan (2009).

If the distribution of the agents’ valuations is known and we assume common knowledge among the agents over those priors, the strategyproofness requirement can be weakened to Bayesian incentive compatibility. In that weaker framework, mechanisms can extract full expected efficiency and achieve budget balance (d’Aspremont and Gérard-Varet, 1979; Arrow, 1979). However, those mechanisms use knowledge of the priors. Therefore, in the general quasi-linear setting, for mechanisms without priors, it is an important open question to characterize the class of strategyproof budget-balanced mechanisms, to find such mechanisms that minimize inefficiency, and to find strategyproof mechanisms that minimize the sum (or other convex combination) of inefficiency and budget imbalance. This paper addresses this important research gap in the general quasi-linear setting, for both deterministic and randomized settings. Our approach is also prior-free—the strategyproofness guarantees consider the worst-case scenarios. We show that the answers are asymptotically positive: even in such a general setup, the Green-Laffont impossibility is not too restrictive when the number of agents is large, and our mechanisms seem to work well on real-world datasets.

2 Model and definitions

We denote the set of agents by $N = \{1, 2, \dots, n\}$ and the set of alternatives by $A = \{a_1, a_2, \dots, a_m\}$. We assume that each agent’s valuation is drawn from an open interval $(-\frac{M}{2}, \frac{M}{2}) \subset \mathbb{R}$, that is, the valuation of agent i is a mapping $v_i : A \rightarrow (-\frac{M}{2}, \frac{M}{2})$, $\forall i \in N$ and is a private information. Denote the set of all such valuations of agent i as V_i and the set of valuation profiles by $V = \times_{i \in N} V_i$.

A *mechanism* is a tuple of two functions $\langle f, \mathbf{p} \rangle$, where f is called the *social choice function* (SCF) that selects the *allocation* and $\mathbf{p} = (p_1, p_2, \dots, p_n)$ is the vector of *payments*, $p_i : V \rightarrow$

$\mathbb{R}, \forall i \in N$. The utility of agent i for an alternative a and valuation profile $v \equiv (v_i, v_{-i})$ is given by the *quasi-linear* function: $v_i(a) - p_i(v_i, v_{-i})$. For *deterministic* mechanisms, $f : V \rightarrow A$ is a deterministic mapping, while for *randomized* mechanisms, the allocation function f is a lottery over the alternatives, that is, $f : V \rightarrow \Delta A$. With a slight abuse of notation, we denote $v_i(f(v_i, v_{-i})) \equiv \mathbb{E}_{a \sim f(v_i, v_{-i})} v_i(a) = \sum_{a \in A} f_a(v_i, v_{-i}) v_i(a)$ to be the expected valuation of agent i for a randomized mechanism. The following definitions are standard in the mechanism design literature.

DEFINITION 1 (Strategyproofness) A mechanism $\langle f, \mathbf{p} \rangle$ is strategyproof if for all $v \equiv (v_i, v_{-i}) \in V$,

$$v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}), \quad \forall v'_i \in V_i, i \in N.$$

DEFINITION 2 (Efficiency) An allocation f is efficient if it maximizes social welfare, that is, $f(v) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} v_i(a)$, $\forall v \in V$.

DEFINITION 3 (Budget Balance) A payment function $p_i : V \rightarrow \mathbb{R}, i \in N$ is budget balanced if $\sum_{i \in N} p_i(v) = 0$, $\forall v \in V$.

In addition, in parts of this paper we will consider mechanisms that are oblivious to the alternatives—a property known as *neutrality*. To define this, we consider a permutation $\pi : A \rightarrow A$ of the alternatives. Therefore, π over a randomized mechanism and over a valuation profile will imply that the probability masses and the valuations of the agents are permuted over the alternatives according to π , respectively.³

DEFINITION 4 (Neutrality) A mechanism $\langle f, \mathbf{p} \rangle$ is neutral if for every permutation of the alternatives π (where $\pi(v) \neq v$) we have

$$\pi(f(v)) = f(\pi(v)) \quad \text{and} \quad p_i(\pi(v)) = p_i(v), \quad \forall v \in V, \forall i \in N.$$

Note that efficient social choice functions are neutral and the Green-Laffont result implicitly assumes this property.

The most important class of allocation functions in the context of deterministic mechanisms are *affine maximizers*, defined as follows.

DEFINITION 5 (Affine Maximizers) An allocation function f is an affine maximizer if there exist real numbers $w_i \geq 0, i \in N$, not all zeros, and a function $\kappa : A \rightarrow \mathbb{R}$ such that $f(v) \in \operatorname{argmax}_{a \in A} (\sum_{i \in N} w_i v_i(a) + \kappa(a))$.

As we will explain in the body of this paper, we will focus on *neutral* affine maximizers (Mishra and Sen, 2012), where the function κ is zero.

$$f(v) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} w_i v_i(a) \quad \text{neutral affine maximizer} \quad (1)$$

The following property of the mechanism ensures that two different payment functions of an agent, say i , that implement the same social choice function differ from each other by a function that does not depend on the valuation of agent i .⁴

³We have overloaded the notation of π following the convention in social choice literature (see, e.g., Myerson (2013)). The notation $\pi(v)$ denotes the valuation profile where the alternatives are permuted according to π .

⁴This definition is a generalization of auction revenue equivalence and is commonly used in the social choice literature (see, e.g., Heydenreich et al. (2009)).

DEFINITION 6 (Revenue Equivalence) *An allocation f satisfies revenue equivalence if for any two payment rules p and p' that make f strategyproof, there exist functions $h_i : V_{-i} \rightarrow \mathbb{R}$, such that*

$$p_i(v_i, v_{-i}) = p'_i(v_i, v_{-i}) + h_i(v_{-i}), \quad \forall v_i \in V_i, \forall v_{-i} \in V_{-i}, \forall i \in N.$$

The metrics of inefficiency we consider in this paper are defined as follows.

DEFINITION 7 (Sample Inefficiency) *The sample inefficiency for a deterministic mechanism $\langle f, \mathbf{p} \rangle$ is:*

$$r_n^M(f) := \frac{1}{nM} \sup_{v \in V} \left[\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right]. \quad (2)$$

The metric is adapted to expected sample inefficiency for randomized mechanisms:

$$r_n^M(f) := \frac{1}{nM} \sup_{v \in V} \left\{ \mathbb{E}_{f(v)} \left[\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right] \right\}. \quad (3)$$

The majority of this paper is devoted to finding strategyproof and budget balanced mechanisms $\langle f, \mathbf{p} \rangle$ that minimize the sample inefficiency.

A different, but commonly used, metric of inefficiency in the literature is the worst-case ratio of the social welfare of the mechanism and the maximum social welfare: $\inf_{v \in V} \frac{\sum_{i \in N} v_i(f(v))}{\max_{a \in A} \sum_{i \in N} v_i(a)}$. A conclusion similar to what we prove in this paper: “*inefficiency vanishes when $n \rightarrow \infty$* ”, holds in that metric as well, but unlike our metric, that metric would require an additional assumption that the valuations are positive, which is not always the case in a quasi-linear domain.

We are now ready to start presenting our results. We begin with deterministic mechanisms that are strategyproof and budget balanced.

3 Deterministic, strategyproof, budget-balanced mechanisms

Before presenting the main result of this section, we formally define a class of mechanisms we call *sink* mechanisms. A sink mechanism has one or more *sink* agents, given by the set $S \subset N$, picked a priori, whose valuations are not used when computing the allocation (i.e., $f(v) = f(v_{-S})$) and the sink agents do not pay anything and together they receive the payments made by the other agents. The advantage of a sink mechanism is that it is strategyproof if it is strategyproof for the agents other than the sink agents and the surplus is divided among the sink agents in some reasonable manner, and sink mechanisms are budget balanced by design. An example of a sink mechanism is where $S = \{i_s\}$ (only one sink agent) and $f(v_{-i_s})$ chooses an alternative that would be efficient had agent i_s not exist, that is, $f(v_{-i_s}) = \operatorname{argmax}_{a \in A} \sum_{i \in N \setminus \{i_s\}} v_i(a)$. The [Clarke \(1971\)](#) payment rule can be applied here to make the mechanism strategyproof for the rest of the agents—that is, for agents other than i_s , $p_i(v_{-i_s}) = \max_{a \in A} \sum_{j \in N \setminus \{i_s, i\}} v_j(a) - \sum_{j \in N \setminus \{i_s, i\}} v_j(f(v_{-i_s}))$, $\forall i \in N \setminus \{i_s\}$. Paying agent i_s the ‘leftover’ money (that is, $p_{i_s}(v_{-i_s}) = -\sum_{j \in N \setminus \{i_s\}} p_j(v_{-i_s})$) makes the mechanism budget balanced. Our first result establishes that the existence of a sink agent is not only sufficient but also *necessary* for deterministic mechanisms.

THEOREM 1 *Any deterministic, strategyproof, budget-balanced, neutral mechanism $\langle f, \mathbf{p} \rangle$ in the domain V has at least one sink agent.*⁵

⁵Green and Laffont’s impossibility result holds for efficient mechanisms, and all efficient mechanisms are neutral. However, the neutrality of an efficient rule is upto tie-breaking, and Green-Laffont applies no matter how the tie is broken. Similarly, our result also holds irrespective of how the tie is broken. Therefore, this theorem covers and generalizes that result since having at least one sink agent implies that the outcome cannot be efficient.

The proof involves two steps.

1. We leverage the fact that a mechanism that satisfies the stated axioms must necessarily be a neutral affine maximizer (Equation (1)) and has a specific structure for payments. The characterization of the payment structure comes from revenue equivalence.
2. The core of the proof then lies in showing that for such payment functions, it is impossible to have no sink agents (identified as agents that have zero weights, $w_i = 0$, in the affine maximizer). This is shown in a contrapositive manner—assuming that there is no sink agent, we construct valuation profiles that lead to a contradiction to budget balance.

The complete proof is given in the appendix.

Our next goal is to find the mechanism in this class that gives the *lowest* sample inefficiency (Equation (2)). In the proof of the next theorem (presented in the appendix) we show that this is achieved when there is exactly one sink and the neutral affine maximizer weights are equal for all agents other agents. This, in turn, yields the following lower bound on inefficiency.

THEOREM 2 *For every deterministic, strategyproof, budget-balanced, neutral mechanism $\langle f, \mathbf{p} \rangle$ over V , $r_n^M(f) \geq \frac{1}{n}$. This bound is tight.*

4 Jointly minimizing budget imbalance and inefficiency

In the previous section, we considered strategyproof, budget-balanced mechanisms that are minimally inefficient. We achieved a sample inefficiency lower bound of $\frac{1}{n}$. *Could one do better by, instead of requiring budget balance and minimizing inefficiency, relaxing budget balance by allowing money burning, and then minimizing the inefficiency from the allocation plus the inefficiency caused by money burning (or required subsidy from outside the mechanism)?* This would seem like the sensible objective for minimizing overall waste.

In this section, we consider the joint problem of minimizing inefficiency and budget imbalance in the setting of deterministic mechanisms. We consider a convex combination of these two quantities since in the quasi-linear domain both of them contribute additively in the agents' utilities and social welfare. We assume that the combination proportions are normative constants with which the planner associates importance to the two factors of inefficiency and budget balance. Therefore, the coefficients of the convex combination are independent of the number of agents. We show that for unrestricted valuations, i.e., when the valuation bound M is large, considering this joint problem does not yield a better than $1/n$ rate of decay of the *efficiency-budget spillover* defined as follows.

$$\rho_n(f, \mathbf{p}) := \lim_{M \rightarrow \infty} \frac{1}{nM} \sup_{v \in V} [\lambda \cdot T_1^n(f, v) + (1 - \lambda) \cdot T_2^n(\mathbf{p}, v)], \quad (4)$$

Where $T_1^n(f, v) = (\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)))$ and $T_2^n(\mathbf{p}, v) = |\sum_{i \in N} p_i(v)|$. For $\lambda = 1$, that is, when budget imbalance is not a concern, one can use the VCG mechanism to get $\rho_n(f, \mathbf{p}) = 0$. Similarly, for $\lambda = 0$, a sink mechanism will give $\rho_n(f, \mathbf{p}) = 0$. So, the interesting cases are when $\lambda \in (0, 1)$, and for this we have a solution that decays as $1/n$. In this section, we will assume that $\lambda, 0 < \lambda < 1$ is exogenous. Our goal is to find a strategyproof and neutral mechanism $\langle f, \mathbf{p} \rangle$ that minimizes the objective ρ_n . We have shown in Section 3 that $T_1^n(f, v)$ can at most be a constant when $T_2^n(\mathbf{p}, v)$ is zero for every v . Hence, for any improvement in the efficiency-budget spillover metric, that is, for $\rho_n(f, \mathbf{p}) = o(r_n(f))$, it is necessary that the

term $\sup_{v \in V} [\lambda \cdot T_1^n(f, v) + (1 - \lambda) \cdot T_2^n(\mathbf{p}, v)]$ be $o(1)$. Since both $T_1^n(f, v)$ and $T_2^n(\mathbf{p}, v)$ are non-negative, it is necessary that the factor $T_2^n(\mathbf{p}, v) = o(1)$ for every $v \in V$. Our next result shows that it is impossible to have $T_2^n(\mathbf{p}, v) = o(1)$, $\forall v \in V \Leftrightarrow \lim_{n \rightarrow \infty} \sup_{v \in V} T_2^n(\mathbf{p}, v) = 0$. Hence, for deterministic mechanisms with unrestricted valuations, the bound on inefficiency *with no budget imbalance* (presented in Section 3) is *asymptotically* optimal for this joint optimization problem as well.⁶

THEOREM 3 (Unimprovability) *For every deterministic, strategyproof, and neutral mechanism $\langle f, \mathbf{p} \rangle$ over V and for every $\lambda \in (0, 1)$, $\rho_n(f, \mathbf{p}) = \Omega(\frac{1}{n})$. This bound is tight. For $\lambda = 0$, a sink mechanism, and for $\lambda = 1$, the VCG mechanism, achieves zero spillover.*

5 Randomized, strategyproof, budget-balanced mechanisms

In Section 3, we saw that the best sample inefficiency achieved by a deterministic budget balanced mechanism is $\frac{1}{n}$. In this section, we discuss how the inefficiency can be reduced by considering randomized mechanisms.

An intuitive approach is to consider a mechanism where each agent is picked as a sink with probability $\frac{1}{n}$.

DEFINITION 8 (Naïve Randomized Sink) *A naïve randomized sink (NRS) mechanism picks every agent as a sink w.p. $\frac{1}{n}$ and takes the efficient allocation without that agent. The payments of the non-sink agents are VCG payments without the sink. The surplus is transferred to the sink.*

Clearly, this mechanism is strategyproof, budget balanced, and neutral by design. One can anticipate that this may not yield the best achievable inefficiency bound. Unlike deterministic mechanisms, very little is known about the structure of randomized strategyproof mechanisms in the general quasi-linear setting. Furthermore, we consider mechanisms that are budget-balanced in addition. Hence, even though we can obtain an upper bound on the expected sample inefficiency ($r_n^M(f)$) by considering specific mechanisms like the NRS mechanism described above, the problem of providing a lower bound (i.e., no randomized mechanism can achieve a smaller $r_n^M(f)$ than a given number), seems elusive in the general quasi-linear setting.

Therefore, in the following two subsections, we consider two approaches, respectively. First, we show lower bounds in a special class of strategyproof, budget-balanced, randomized mechanisms. Second, we analytically provide a lower bound of the optimal, strategyproof, budget-balanced, randomized mechanism for two agents and two alternatives, using a discrete relaxation of the original problem (in the spirit of *automated mechanism design* (Conitzer and Sandholm, 2002; Sandholm, 2003)). This approach provides an approximation to the strategyproofness and optimality of the original problem. However, the problems of finding a mechanism that matches this lower bound and extending the lower bound to any number of agents and alternatives are left as future work.

5.1 Generalized sink mechanisms

In the first approach, we consider a broad class of randomized, budget-balanced mechanisms, which we coin *generalized sink mechanisms*. In this class, the probability of an agent i to become a sink is dependent on the valuation profile $v \in V$, and we consider mechanisms with only *one* sink, i.e., if the probability vector returned by a generalized sink mechanism is $g(v)$, then w.p. $g_i(v)$, agent i is treated as the *only* sink agent. (One can think of a more general class of sink mechanisms

⁶It is easy to see that the conclusions of Footnote 5 and theorem 2 hold even under the assumption of large M .

where multiple agents are treated as sink agents simultaneously. However, it is easy to see—by a similar argument to that in the context of deterministic mechanisms, just before Lemma 2—that using multiple sinks cannot decrease inefficiency.) Clearly, the naïve randomized sink mechanism belongs to this class. Once agent i is picked as a sink, the alternative chosen is the *efficient* one *without* agent i . All agents $j \neq i$ are charged a Clarke tax payment in the world without i , and the surplus amount of money is transferred to the sink agent i . Algorithm 1 shows the steps of a generic mechanism in this class.

Algorithm 1 Generalized Sink Mechanisms, \mathcal{G}

- 1: **Input:** a valuation profile $v \in V$
 - 2: A generic mechanism in this class is characterized by a probability distribution over the agents N (which may depend on the valuation profile), $g : V \rightarrow \Delta N$
 - 3: The mechanism randomly picks one agent i in N with probability $g_i(v)$
 - 4: Treat agent i as the sink
-

Clearly, not every mechanism in this class is strategyproof. The crucial aspect is how the probabilities of choosing the sink are decided. If the probability $g_i(v)$ depends on the valuation of agent i , that is, v_i , then there is a chance for agent i to misreport v_i to have higher (or lower) probability of being a sink (being a sink could be beneficial since she gets all the surplus). For example, the *irrelevant sink* mechanism given in Algorithm 2 is *not* strategyproof.

Algorithm 2 Irrelevant Sink Mechanism (not strategyproof)

- 1: **Input:** a valuation profile $v \in V$
 - 2: **for** agent i in N **do**
 - 3: Define: $a^*(v_{-i}) = \operatorname{argmax}_{a \in A} \sum_{j \neq i} v_j(a)$
 - 4: **if** $\sum_{j \neq i} v_j(a^*(v_{-i})) - \sum_{j \neq i} v_j(a) > M$ for all $a \in A \setminus \{a^*(v_{-i})\}$ **then**
 - 5: Call i an irrelevant agent
 - 6: **if** irrelevant agent is found **then**
 - 7: Arbitrarily pick one of them as a sink with probability 1
 - 8: **else**
 - 9: Pick an agent i with probability $\frac{1}{n}$ and treat as sink
-

The intuition of this mechanism is that if agent i 's maximum valuation sweep ($-M/2$ to $M/2$) cannot change the alternative, this *irrelevant* agent can be selected as a sink, which yields the efficient alternative. However, when there is no such irrelevant agent, the decision of choosing every agent equi-probably leads to a chance of manipulation. An agent whose true valuation report does not lead her to become a sink can misreport a valuation so that there is no irrelevant agent, thereby increasing her own probability of being selected as a sink. For example, consider two valuation profiles with three agents (numbered 1, 2, 3) and three alternatives (a, b, c) , and $M = 1$. The agents' valuations in the first profile are $v_1 = (0.5, 0, -0.5)$, $v_2 = (-0.5, 0, 0.5)$, $v_3 = (0, -0.5, 0.5)$, and in the second profile they are $v'_1 = (0.5, 0, -0.5)$, $v'_2 = (-0.5, 0, 0.5)$, $v'_3 = (-0.5, 0, 0.5)$. The mechanism returns agent 1 as the irrelevant agent in the first profile and therefore picks alternative c with probability 1. There is no irrelevant agent in the second profile and hence each agent is picked as a sink with uniform probability, leading to the probability vector $(2/3, 0, 1/3)$ for the

alternatives a, b, c . But agent 3 strictly gains by moving from the first profile to the second.⁷

A small modification of the previous mechanism leads to a strategyproof generalized sink mechanism. This shows that the class of generalized sink mechanisms is indeed richer than the constant probability sink mechanisms. In the modified version, we pick a default sink with a certain probability, which will be the sink if there exists no irrelevant agent among the rest of the agents. The change here is that when an agent is picked as a default sink, her valuation has no effect in deciding the sink. See Algorithm 3.

Algorithm 3 Modified Irrelevant Sink Mechanism (strategyproof)

- 1: **Input:** a valuation profile $v \in V$
 - 2: Pick agent i as a *default sink* with probability p_i
 - 3: **for** agent j in $N \setminus \{i\}$ **do**
 - 4: **if** irrelevant agent(s) found within $N \setminus \{i\}$ **then**
 - 5: Arbitrarily pick one of them as a sink
 - 6: Irrelevant agent is found
 - 7: **if** no irrelevant agent is found within $N \setminus \{i\}$ **then**
 - 8: Treat agent i as sink
-

It is easy to verify that this mechanism is strategyproof. Interestingly, no generalized sink mechanism can improve the expected sample inefficiency over deterministic mechanisms if there are more alternatives than agents ($m > n$).

THEOREM 4 (Generalized Sink for $m > n$) *If $m > n$, every generalized sink mechanism has expected sample inefficiency $\geq \frac{1}{n}$.*

The proof is critically dependent on $m > n$. However, we can hope for a smaller inefficiency if the number of alternatives is small. We state this intuition formally as follows.

THEOREM 5 (Increasing Inefficiency with m) *For every mechanism f and for a fixed number of agents n , the expected sample inefficiency is non-decreasing in m , i.e., $r_{n,m_1}^M(f) \geq r_{n,m_2}^M(f), \forall m_1 > m_2$.⁸*

Theorems 4 and 5 suggest that in order to minimize inefficiency, one must have a small number of alternatives. So from now on, we consider the extreme case with $m = 2$, where we investigate the advantages of randomization.

For two alternatives, the following theorem shows that the naïve randomized sink (NRS) mechanism reduces the inefficiency by a factor of two.

THEOREM 6 (Naïve Randomized Sink) *For $m = 2$, the expected sample inefficiency of the NRS mechanism is $\frac{1}{n^2} \lceil \frac{n}{2} \rceil \sim \frac{1}{2n}$.*

Even though the modified irrelevant sink (MIS) mechanism (Algorithm 3) is sophisticated in its use of the valuation profile, it is easy to check that even that mechanism yields the same inefficiency on the profile illustrated in the proof above.

⁷One can also verify that the weak monotonicity condition, which is a necessary condition for strategyproofness, is violated for agent 3 between these two profiles.

⁸We overload the notation for the expected sample inefficiency r_n with $r_{n,m}$ to make the number of alternatives explicit.

THEOREM 7 (Modified Irrelevant Sink) *For $m = 2$, the expected sample inefficiency of the MIS mechanism (Algorithm 3) is at least $\frac{1}{n^2} \lceil \frac{n}{2} \rceil \sim \frac{1}{2n}$.*

Furthermore, it turns out that NRS and MIS have the same inefficiency on every valuation profile. Both mechanisms choose a single agent as a sink. The default sink for MIS is chosen uniformly at random, identical to the choice of the sink for NRS. If there does not exist an irrelevant sink in the rest of the agents, the inefficiency remains the same as that for the default sink, which is identical to the inefficiency of NRS for that choice of sink. But even if an irrelevant sink exists, by the construction of the irrelevant sink, the resulting alternative is the efficient alternative for the agents except the default sink. This outcome would have resulted even if the default sink was chosen as the sink. Therefore, the inefficiencies in MIS and NRS mechanisms are the same.

The above result does not say much about the lowest achievable expected sample inefficiency (even in this special class of generalized sink mechanisms). In order to understand the limit of lowest achievable inefficiency for randomized mechanisms, we take an optimization-based approach. However, to keep the analysis simple and tractable, we focus on the special case of two agents and two alternatives. Our next result gives a lower bound on the inefficiency for the class of generalized sink mechanisms in that setting. Since we now fix the number of agents in the analysis, minimizing the expected sample inefficiency is equivalent to minimizing the expected *absolute* inefficiency given by $nr_n^M(f)$ which is $\frac{1}{M} \sup_{v \in V} \{ \mathbb{E}_{f(v)} [\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v))] \}$. Without loss of generality we will assume $M = 1$. For the rest of this section, we let ‘inefficiency’ mean the expected absolute inefficiency.

THEOREM 8 (Lower Bound of Generalized Sink) *For $n = m = 2$, the expected absolute inefficiency of every strategyproof generalized sink mechanism is lower bounded by $\frac{1}{2}$.*

5.2 Unrestricted randomized mechanisms

We now move on to study optimal randomized mechanisms without restricting attention necessarily to generalized sink mechanisms. Finding a mechanism that achieves the minimum absolute inefficiency can be posed as the following optimization problem.

$$\begin{aligned}
\min_{f, \mathbf{p}} \quad & \sup_{v \in V} \left[\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right] \\
\text{s.t.} \quad & v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \\
& \geq v_i(f(v'_i, v_{-i})) - p_i(v'_i, v_{-i}), \quad \forall v_i, v'_i, v_{-i}, \forall i \in N \\
& \sum_{a \in A} f_a(v) = 1, \quad \forall v \in V, \\
& \sum_{i \in N} p_i(v) = 0, \quad \forall v \in V, \\
& f_a(v) \geq 0, \quad \forall v \in V, a \in A.
\end{aligned} \tag{5}$$

The objective function denotes the absolute inefficiency. The first set of inequalities in the constraints denote the strategyproofness requirement, where the term $v_i(f(v)) = v_i \cdot f(v)$ denotes the expected valuation of agent i due to the randomized mechanism f . The second and last set of inequalities ensure that the $f_a(v)$ ’s are valid probability distributions, and the third set of inequalities ensure that the budget is balanced. The optimization is over the social choice functions f

and the payments \mathbf{p} , where the f variables are non-negative but the p variables are unrestricted. Clearly, this is a linear program (LP), which has an uncountable number of constraints (because the equalities and inequalities have to be satisfied at all $v \in V$, which are the profiles of valuation functions mapping alternatives to an open interval). We address this optimization problem using finite constrained optimization techniques by discretizing the valuation levels. We assume that each agent's valuations are uniformly discretized with k levels in $[-M/2, M/2]$, which makes the set of valuation profiles V finite. The optimal value of such a discretized relaxation of the constraints provides a lower bound on the optimal value of the original problem. This is because the discretized relaxation of the valuations only increases the feasible set since some of the constraints are removed, that is, more f 's and p 's satisfy the constraints, allowing a potentially lower value to be achieved for the minimization objective. We now prove a lower bound when the number of discretized levels is three. The analysis uses a primal-dual argument on the discrete relaxation of the problem.

THEOREM 9 (Lower Bound of Inefficiency for Randomized Mechanisms) *For $n = m = 2$, and for $k = 3$ discrete levels of valuations, the absolute inefficiency is lower bounded by $\frac{1}{7} = 0.142857$.*




Figure 1: Lower bound for the discrete relaxation of the inefficiency minimization LP.

The proof technique can be extended to a larger number of discrete levels to obtain a tighter lower bound on the actual inefficiency. We conducted a form of automated mechanism design (Conitzer and Sandholm, 2002; Sandholm, 2003) by solving this LP using Gurobi (2015) for increasing values of k . We apply the same optimization-based approach for generalized sink and the deterministic cases as well, even though for these cases we have theoretical bounds. The solid lines in Figure 1 show the optimization-based results (denoted as AMD) and the dotted lines show the theoretical bounds. Note that for deterministic case, the theoretical and optimization-based approaches overlap since the inefficiency is unity even with two valuation levels. The convergence of the optimization-based approach for generalized sink mechanism shows the efficacy of the approach and helps to predict the convergence point for the optimal randomized mechanism. One can see

that the lower bound is greater than 0.2 for the optimal mechanism, but it seems to converge to a value much lower than 0.5.

6 Experiments with real data

Even though the naïve randomized sink (NRS) (Definition 8) mechanism gives a worst-case sample inefficiency between $1/2n$ (for 2 alternatives, Theorem 6) and $1/n$ (for more alternatives than agents, $m > n$, Theorem 4), in this section we investigate its average and worst-case performances on real datasets of user preferences. Going back to the example of movie selection by a group of friends (Section 1), we consider several sizes of the group. A small group consists of tens of friends, while if the decision involves screening a movie at a school auditorium, the group size could easily be in the hundreds. This is why we consider group sizes spanning from 10 to 210 in steps of 50.

A similar situation occurs when a group of people decides which comedian/musician to invite in a social gathering, where they need to pay the cost of bringing the performer.

Keeping these motivating situations in mind, we used two datasets that closely represent the scenarios discussed. We used the MovieLens 20M dataset (Harper and Konstan, 2016) and the Jester dataset (Goldberg et al., 2001) to compare the performance of the two mechanisms with their worst case bounds. The first dataset contains preferences for movies, while the second contains preferences for online jokes. The MovieLens 20M dataset (ml-20m) describes users’ ratings between 1 and 5 stars from MovieLens, a movie recommendation service. It contains 20,000,263 ratings across 27,278 movies. These data were created from the ratings of 138,493 users between January 09, 1995 and March 31, 2015. For our experiment, we sampled the preferences of a specific number of users (shown as agents on the x-axis of Figure 2) multiple times uniformly at random from the whole set of users that rated a particular genre of movies, and computed the sample inefficiency on this sampled set and plotted the average expected sample inefficiency and standard deviation.

The Jester dataset (jester-data-1) used in our experiment contains data from 24,983 users who have rated 36 or more jokes, a matrix with dimensions 24983 X 100, and is obtained from Jester, an online joke recommendation system.⁹

Figure 2a shows that the real preferences of users yield much lower expected sample inefficiencies for the naïve randomized sink (NRS) mechanism than the theoretical worst-case guarantee. The improvement ranges from roughly a factor of 5 (for a group size of 10) to almost 100 (for a group size of 210). This also indicates that the rate of decay of the inefficiency with the size of the group is faster than the theoretical guarantee. The experiment is run for different sizes of the agent groups with a random group of each size drawn multiple times from the dataset. The bars in Figure 2a shows the *average* expected sample inefficiencies of the mechanism with the standard deviations around them. Since NRS is a randomized mechanism, it is also worth looking at its worst-case performance on the same sampled datasets. Figure 2b shows the *average* worst-case sample inefficiency along with its standard deviation. The line plot in both the figures shows the worst-case expected sample inefficiency for NRS for two alternatives.

By the arguments following Theorem 7, we know that the inefficiency of the modified irrelevant sink (MIS) (Algorithm 3) will be same as NRS. Since the MIS mechanism also picks exactly one sink, the worst-case behavior of a single sink illustrated above also applies to the MIS mechanism.

⁹In both datasets there are missing values because a user has typically not rated all movies/jokes. Before our experiment, we filled the missing values with a random realization of ratings drawn from the empirical distribution for that alternative (movie or joke). The empirical distribution of an alternative is created from the histogram of the available ratings of the users. We cleaned the dataset by keeping only those alternatives that have at least 10 or more available ratings and filled the rest using their empirical distributions.

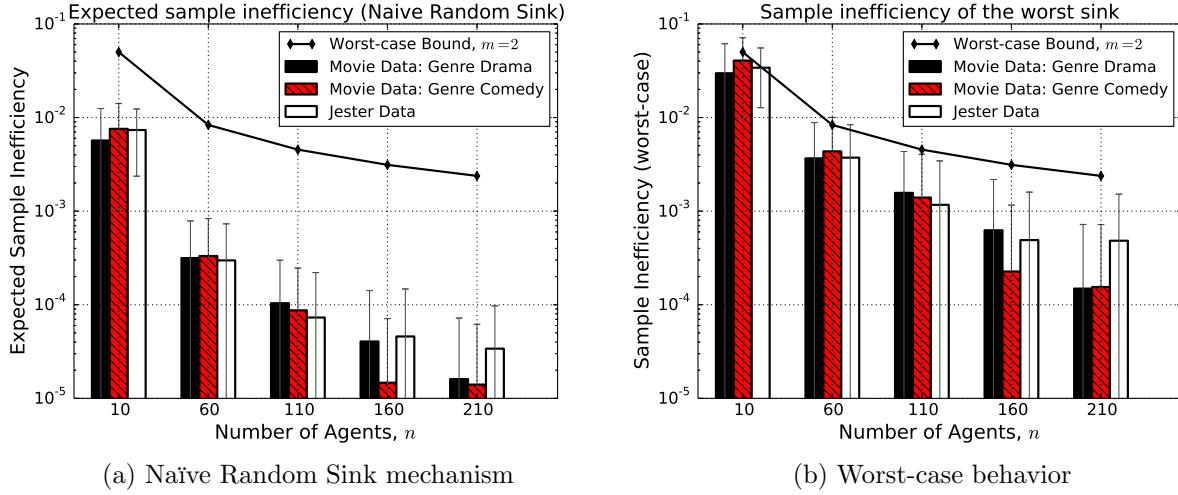


Figure 2: Performance of naïve randomized sink mechanism on MovieLens and Jester datasets.

7 Conclusions and future research

In this paper, we provided several new results on the classic question of the interplay between efficiency and budget balance, properties that are incompatible with strategyproofness due to the Green-Laffont impossibility result, in the general quasi-linear setting. We sought to understand the limits of minimal compromise between these two properties, both in the context of deterministic and randomized mechanism design framework.

We proved characterization results, and a tight lower bound for inefficiency, for deterministic budget-balanced mechanisms. We also proved that for unrestricted valuations, minimizing inefficiency and budget imbalance together does not provide any asymptotic advantage in the deterministic paradigm over requiring budget balance and minimizing inefficiency.

We proved that randomization helps—particularly when the number of alternatives is small compared to the number of agents. Motivated by our result for deterministic mechanisms that shows that a strategyproof, budget-balanced, neutral mechanism must include a sink, we introduced the class of *generalized sink mechanisms* which is a general (and adaptive to the valuation profile) way of picking sink agents. We showed that there exists strategyproof non-trivial mechanism (modified irrelevant sink) in this class that reduces the worst-case inefficiency by a factor of 2. We used an automated mechanism design approach for two agents and showed analytically that an optimal randomized mechanism offers further reduction in the inefficiency.

Experiments with real data from two applications compare the naïve randomized sink with its theoretical worst-case upper bound. We see that the mechanism perform well in practice and yield very little inefficiency ($\sim 1\%$ to 0.01% depending on the group size). This inefficiency is 5–100 times smaller than the worst case, and this relative difference increases with the number of agents. We also consider the worst-case realization of this mechanism, and found that the sample inefficiency is close to the worst-case expected sample inefficiency of the mechanism.

Future research includes studying the structure of the optimal randomized mechanisms that achieve the (theoretical) improved efficiency. Future work also includes investigating the rate of improvement of the optimal bound for a general number of agents.

References

- Kenneth Arrow. *The property rights doctrine and demand revelation under incomplete information*. Economics and Human Welfare. New York Academic Press, 1979.
- Itai Ashlagi, Felix Fischer, Ian A Kash, and Ariel D Procaccia. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*, pages 1–13, 2013.
- Martin J Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91(2): 107–126, 1997.
- Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the conference on autonomous agents and multiagent systems (AAMAS)*, pages 882–889, 2006.
- Edward Clarke. Multipart Pricing of Public Goods. *Public Choice*, (8):19–33, 1971.
- Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110. Morgan Kaufmann Publishers Inc., 2002.
- C. d’Aspremont and L. A. Gérard-Varet. Incentives and Incomplete Information. *Journal of Public Economics*, 11(1):25–45, 1979.
- Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In *Proceedings of the annual ACM Symposium on Theory of Computing*, pages 644–652. ACM, 2006.
- Boi Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce VI. Theories for and Engineering of Distributed Mechanisms and Systems*, pages 30–43. Springer, 2004.
- Michal Feldman and Yoav Wilf. Randomized strategyproof mechanisms for facility location and the mini-sum-of-squares objective. *arXiv preprint arXiv:1108.1762*, 2011.
- Robert J Gary-Bobo and Touria Jaaidane. Polling mechanisms and the demand revelation problem. *Journal of Public Economics*, 76(2):203–238, 2000.
- Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins. Eigentaste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4(2):133–151, 2001.
- Jerry R Green and Jean-Jacques Laffont. *Incentives in public decision making*. North-Holland, 1979.
- Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–31, July 1973.
- Sujit P Gujar and Yadati Narahari. Redistribution mechanisms for assignment of heterogeneous objects. *Journal of Artificial Intelligence Research*, pages 131–154, 2011.

- Mingyu Guo and Vincent Conitzer. Optimal-in-expectation redistribution mechanisms. In *Proceedings of the conference on autonomous agents and multiagent systems (AAMAS)*, pages 1047–1054, 2008.
- Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of vcg payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009.
- Mingyu Guo and Vincent Conitzer. Better redistribution with inefficient allocation in multi-unit auctions. *Artificial Intelligence*, 216:287–308, 2014.
- Gurobi. Gurobi optimizer reference manual, 2015. URL <http://www.gurobi.com>.
- F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 5(4):19, 2016.
- Birgit Heydenreich, Rudolf Müller, Marc Uetz, and Rakesh V Vohra. Characterization of revenue equivalence. *Econometrica*, 77(1):307–316, 2009.
- Vijay Krishna and Eliot Maenner. Convex potentials with an application to mechanism design. *Econometrica*, 69(4):1113–1119, 2001.
- Jordi Massó, Antonio Nicolò, Arunava Sen, Tridib Sharma, and Levent Ülkü. On cost sharing in the provision of a binary and excludable public good. *Journal of Economic Theory*, 155:30–49, 2015.
- Debasis Mishra and Arunava Sen. Roberts’ theorem with neutrality: A social welfare ordering approach. *Games and Economic Behavior*, 75(1):283–298, 2012.
- Debasis Mishra and Tridib Sharma. On the optimality of the green-laffont mechanism. Preliminary Draft, 2016.
- Hervé Moulin. Almost budget-balanced vcg mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.
- Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
- Roger B Myerson. Fundamentals of social choice theory. *Quarterly Journal of Political Science*, 8(3):305–337, 2013.
- Ariel D Procaccia and Moshe Tennenholtz. Approximate mechanism design without money. In *Proceedings of the 10th ACM conference on Electronic Commerce (EC)*, pages 177–186. ACM, 2009.
- Ralph Tyrrell Rockafellar. *Convex analysis*, volume 28. Princeton university press, 1997.
- Michael H. Rothkopf. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research*, 55(2):191–197, 2007.
- Tim Roughgarden and Mukund Sundararajan. Quantifying inefficiency in cost-sharing mechanisms. *Journal of the ACM (JACM)*, 56(4):23, 2009.

Tuomas Sandholm. Automated mechanism design: A new application area for search algorithms. In *Principles and Practice of Constraint Programming–CP 2003*, pages 19–36. Springer, 2003.

Nguyen Kim Thang. On randomized strategy-proof mechanisms without payment for facility location games. In *Proceedings of the Web and Internet Economics (WINE). Stanford, USA*, pages 13–16, 2010.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

Appendix

Proof of Footnote 5

Proof: Consider the class of deterministic, strategyproof, and neutral mechanisms. Mishra and Sen (2012) have shown that in the domain V , an allocation that satisfies the properties above must be a neutral affine maximizer (Definition 5), that is, there exists $w_i \geq 0, \forall i \in N$, not all zero, such that,

$$f(v) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} w_i v_i(a). \quad (6)$$

Additionally, the result by Rockafellar (1997) and Krishna and Maenner (2001) states that for any convex type space, if the valuations are linear in type, then a strategyproof allocation satisfies revenue equivalence (Definition 6). In our setting, the types of the agents are their valuations, which implies, trivially, that the valuations are linear in type. Also, they are drawn from the interval $(-\frac{M}{2}, \frac{M}{2})$, which is convex. So, revenue equivalence holds for the allocations in our setting. The following payment implements the affine maximizer allocation f given by Equation (6):

$$p_i(v_i, v_{-i}) = \begin{cases} -\frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right), & w_i > 0 \\ 0, & w_i = 0 \end{cases} \quad (7)$$

for all $i \in N$. Since revenue equivalence holds in this setting, we conclude that any payment $\hat{p}_i, i \in N$ that makes $\langle f, \hat{\mathbf{p}} \rangle$ strategyproof, will be different from the above mentioned payments \mathbf{p} by an additive factor $h_i(v_{-i})$ for each agent i in every valuation profile.

Now, we turn to proving the result of the theorem. We have the functional form of deterministic, strategyproof, neutral mechanisms given by Equation (6). If, on this class of mechanisms, we show that one cannot have weights $w_i > 0$ for all $i \in N$ while imposing budget balance, then we are done. This is because, if there exists one agent $i \in N$, for which $w_i = 0$, that agent is a *sink* agent as her valuations are never used by the social choice function and she is charged no payment. By revenue equivalence, any other payment that can implement the same allocation f is $h_i(v_{-i})$. Putting this in the budget balance equation, we get $h_i(v_{-i}) = -\sum_{j \in N \setminus \{i\}} p_j(v)$, that is, she receives the payments made by the other agents. Thus agent i is a sink agent. Hence, the proof is completed by proving the following claim.

LEMMA 1 (Existence of $w_i = 0$ Agent) *A budget balanced mechanism $\langle f, \mathbf{p} \rangle$, where f is a neutral affine maximizer on the domain V , must have at least one agent i that has $w_i = 0$.*

Proof: Suppose for contradiction that $w_i > 0, \forall i \in N$. Since f is a neutral affine maximizer (Equation (6)) and revenue equivalence holds in V (Equation (7)), we know that the payments are of the form $p_i(v_i, v_{-i}) = h_i(v_{-i}) - \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right), \forall v \in V, \forall i \in N$.

Additionally, since the mechanism $\langle f, \mathbf{p} \rangle$ is also budget balanced, we have

$$\begin{aligned} & \sum_{i=1}^n \left(h_i(v_{-i}) - \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right) \right) = 0, \forall v \in V \\ \Rightarrow & \sum_{i=1}^n h_i(v_{-i}) - \sum_{i=1}^n \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i(f(v)) = 0, \forall v \in V. \end{aligned} \quad (8)$$

For an easier exposition, we first explain the proof technique for $n = 2$. Later the same proof is generalized to any number of agents.

By assumption, $w_1, w_2 > 0$. Pick two valuation profiles (v_1^+, v_2) and (v_1^-, v_2) such that the affine maximizer alternative in the first is a_1 while that in the second is a_2 , that is,

$$w_1 v_1^+(a_1) + w_2 v_2(a_1) > w_1 v_1^+(a_2) + w_2 v_2(a_2) \quad (9)$$

$$w_1 v_1^-(a_1) + w_2 v_2(a_1) < w_1 v_1^-(a_2) + w_2 v_2(a_2) \quad (10)$$

This can be done by choosing $v_1^+(a_2) = v_1^-(a_2) = v_1(a_2)$ (say) small and v_2 to be small enough for both alternatives, so that the valuation of agent 1 for a_1 determines the resulting alternative of f . Therefore, the RHS of the inequalities above are the same. Since the inequality of Equation (10) is strict, let the difference of the RHS and LHS be $\delta > 0$. The allocations at these two profiles are: $f(v_1^+, v_2) = a_1$ and $f(v_1^-, v_2) = a_2$. Since the payments satisfy revenue equivalence and budget balance, Equation (8) holds, which gives

$$\begin{aligned} & -\frac{1}{w_1} w_2 v_2(a_1) + h_1(v_2) - \frac{1}{w_2} w_1 v_1^+(a_1) + h_2(v_1^+) = 0 \\ & -\frac{1}{w_1} w_2 v_2(a_2) + h_1(v_2) - \frac{1}{w_2} w_1 v_1^-(a_2) + h_2(v_1^-) = 0. \end{aligned}$$

Subtracting the first equation from the second and rearranging, we get

$$\frac{1}{w_1} w_2 (v_2(a_1) - v_2(a_2)) = \frac{1}{w_2} w_1 (v_1^-(a_2) - v_1^+(a_1)) - h_2(v_1^-) + h_2(v_1^+). \quad (11)$$

Note that the RHS is independent of v_2 . Therefore, if $v_2(a_1)$ is increased by a small amount ($< \delta/w_2$), both the inequalities given by Equations (9) and (10) still hold, but Equation (11) fails to hold, which is a contradiction.

The general proof of this lemma extends this idea to any number of agents $n \geq 2$. We prove this for a set of alternatives $A = \{0, 1\}$. Consider this setting as that of a public project. In alternative 0, the project is not undertaken, yielding every agent a value of zero, and when 1 is chosen—i.e., the project is undertaken—the valuation of each agent is denoted by a single real number. This assumption helps us reduce the notational complexity. The proof, however, is completely general for any number of alternatives.

Let the agents be numbered in decreasing order of their weights WLOG, that is, $w_i \geq w_{i+1}, i = 1, 2, \dots, n-1$. We consider the following valuation profile: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1} + \delta, v_n), \delta > 0$ such that

$$-\delta \sum_{i=1}^{n-1} w_i < \sum_{i=1}^n w_i v_i < -\delta \sum_{i=1}^{n-2} w_i \quad (12)$$

The above inequalities imply that the affine maximizer alternative given by Equation (6) for the profile mentioned above is 1. However, if any agent i 's, $i = 1, 2, \dots, n-1$, valuation changes from $v_i + \delta$ to v_i , the alternative changes to 0. We use a generic notation v^S to denote this profile, where S denotes the set of agents such that the valuations of all the agents $k \in S$ are v_k , and the valuations of the agents $j \notin S$ are $v_j + \delta$. Hence, $v^{\{n\}}$ is the profile mentioned before: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1} + \delta, v_n)$ and $v^{\{n-1, n\}}$ is the profile: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1}, v_n)$, for example.

Since, $f(v^{\{n\}}) = 1$, from Equation (8) we have,

$$\left(\sum_{i=1}^{n-1} h_i(v_{-i}^{\{n\}}) + h_n(v_{-n}^{\{n\}}) \right) - \left(\sum_{i=1}^n \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) = 0. \quad (13)$$

The idea of the proof is to make a series of substitutions in the first parentheses of the expression above, leaving the terms in the other parentheses unchanged. Note that, the expression in the second parentheses depends on v_n , while the expression $h_n(v_{-n}^{\{n\}})$ does not. The substitutions sequentially eliminate the dependency on v_n from all the terms in the first parentheses, similar to what we did in the two agent case before. This leads to a contradiction, since v_n can be perturbed arbitrarily small so that it continues to satisfy the inequalities of Equation (12), our only assumed condition, but violates the equality in Equation (13).

The substitutions will involve the term $\sum_{i=1}^{n-1} h_i(v_{-i}^{\{n\}})$ in the first parentheses of Equation (13). Consider the profiles $v^{\{j, n\}}, j = 1, \dots, n-1$. In each of these profiles, $f(v^{\{j, n\}}) = 0$ (due to the choice of $v^{\{n\}}$ in Equation (12)). Hence,

$$\sum_{i=1}^{n-1} h_i(v_{-i}^{\{j, n\}}) + h_n(v_{-n}^{\{j, n\}}) = 0, \quad \forall j \in \{1, \dots, n-1\}. \quad (14)$$

Note that $v_{-i}^{\{i, n\}} = v_{-i}^{\{n\}}$. Hence, we can substitute terms from Equation (14) to the terms in the first parentheses of Equation (13) to get,

$$\left(- \sum_{i=1}^{n-1} \sum_{j \neq \{i, n\}} h_j(v_{-j}^{\{i, n\}}) - \sum_{j \neq n} h_n(v_{-n}^{\{j, n\}}) + h_n(v_{-n}^{\{n\}}) \right) - \left(\sum_{i=1}^n \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) = 0. \quad (15)$$

We continue replacing the terms $h_j(v_{-j}^{\{i, n\}})$ in the first summation of the first parentheses above. All other terms in that parentheses are h_n functions and, therefore, are independent of v_n . For every $i \neq n$, consider the valuation profiles $v^{\{j, i, n\}}, j \neq i, n$. By Equation (12), $f(v^{\{j, i, n\}}) = 0$, so we get an equality similar to Equation (14): $\sum_{k=1}^{n-1} h_k(v_{-k}^{\{j, i, n\}}) + h_n(v_{-n}^{\{j, i, n\}}) = 0, \forall j \neq i, n$. Also, $v_{-j}^{\{j, i, n\}} = v_{-j}^{\{i, n\}}$. So, we follow the same procedure to replace the terms $h_j(v_{-j}^{\{i, n\}})$ in Equation (15)

to yield a similar equality where more terms that were dependent on v_n are now replaced with h_n functions, which are independent of v_n . Since the number of agents is finite, this process will stop after a finite number of iterations, reducing the terms in the first parentheses consisting only of h_n functions. This construction shows that a small perturbation of v_n , which keeps Equation (12) unaffected, will violate the equality obtained through the iterative procedure described above. This completes the proof of the lemma. \blacksquare

Lemma 1 shows that there exists an agent with weight zero, which is a sink agent, and hence the proof of Theorem 1 is complete. \blacksquare

Proof of Theorem 2

Proof: From Footnote 5, we know that any f that satisfies the properties mentioned in the statement of the current theorem must be a neutral affine maximizer with at least one agent i^* that has $w_{i^*} = 0$. We now show that the minimum sample inefficiency $r_n^M(f)$ is achieved when there is exactly *one* such agent i^* and the weights of the other agents $i \in N \setminus \{i^*\}$ are equal. We prove this in two steps: (a) first we show that the proposed single-sink mechanism indeed gives a worst-case sample inefficiency of $1/n$, and having multiple sinks can only make the inefficiency worse, (b) having unequal weights for the rest of the agents is suboptimal for sample inefficiency. This immediately proves the theorem since the proposed mechanism has sample inefficiency $\frac{1}{n}$.

Step (a): the proposed mechanism, having exactly one i^* such that $w_{i^*} = 0$, picks the welfare maximizing allocation not considering the sink agent i^* , that is, $f(v) \in \operatorname{argmax}_{a \in A} \sum_{j \in N \setminus \{i^*\}} v_j(a)$. Denoting $a^*(v)$ to be the efficient allocation, we write:

$$\begin{aligned} \max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) &= \sum_{i \in N} v_i(a^*(v)) - \sum_{i \in N} v_i(f(v)) \\ &= v_{i^*}(a^*(v)) - v_{i^*}(f(v)) + \left[\sum_{j \in N \setminus \{i^*\}} v_j(a^*(v)) - \sum_{j \in N \setminus \{i^*\}} v_j(f(v)) \right] < \left(\frac{M}{2} - \left(-\frac{M}{2} \right) \right) + 0 = M \end{aligned} \quad (16)$$

The first part of the above inequality comes from the fact that the valuations are drawn from $(-\frac{M}{2}, \frac{M}{2})$ so the difference in valuation can at most be M . The second part of the inequality holds because $f(v)$ is the welfare maximizing allocation excluding agent i^* .

It is easy to verify that this inequality is tight at the following valuation profile: $v_{i^*}(a) = \frac{M}{2} - \delta$, $v_{i^*}(z) = -\frac{M}{2} + \gamma$, $\forall z \neq a$, and $v_j(b) = -\frac{M}{2} + \epsilon$, $v_j(z) = -\frac{M}{2} + \frac{\epsilon}{2}$, $\forall z \neq b, \forall j \neq i^*$, where $\delta, \gamma, \epsilon > 0$ are arbitrarily small. The alternatives are: $a^*(v) = a, f(v) = b$. Clearly, this satisfies the above inequality and by taking $\delta, \gamma, \epsilon \rightarrow 0$, we get that the supremum of the difference term approaches M , and hence the sample inefficiency becomes $\frac{1}{n}$.

This counterexample also shows that having more than one sink agent will make the sample inefficiency worse—that is, larger. This is because we can replicate the valuation of i^* for every other sink and the inequality above will be tightly upper bounded at $2M$ for 2 sinks, $3M$ for 3 sinks, etc. Consequently, the sample inefficiency increases to $\frac{2}{n}, \frac{3}{n}$ etc.

To prove step (b), we use the following lemma.

LEMMA 2 (Lowest Sample Inefficiency) *In the class of neutral affine maximizers given by Equation (6) having a sink agent i^* (i.e., $w_{i^*} = 0$), the lowest sample inefficiency is achieved when $w_i = w$ for all $i \in N \setminus \{i^*\}$.*

Proof: Suppose not, that is, $\exists j, j' \in N \setminus \{i^*\}$ such that, WLOG, $w_j > w_{j'}$. Consider the following valuation profile:

$$\begin{aligned} v_{i^*}(a) &= -\frac{M}{2} + \gamma, & v_{i^*}(b) &= \frac{M}{2} - \delta, \\ v_j(a) &= \frac{M}{2} - \delta, & v_j(b) &= \frac{M}{2} - \frac{w_{j'}}{w_j}M - \epsilon, & v_i(a) &= v_i(b), \forall i \in N \setminus \{i^*, j, j'\}, \\ v_{j'}(a) &= -\frac{M}{2} + \gamma, & v_{j'}(b) &= \frac{M}{2} - \delta, & v_i(z) &= -\frac{M}{2} + \frac{\gamma}{2}, \forall z \in A \setminus \{a, b\}, \forall i \in N. \end{aligned}$$

The constants $\delta, \gamma, \epsilon > 0$ are arbitrarily small. It is easy to verify that on this profile, by choosing the constants δ, γ, ϵ appropriately small, the affine maximizer will return a . But the efficient alternative is $a^*(v) = b$. Consider the term $\sum_{i \in N} v_i(a^*(v)) - \sum_{i \in N} v_i(f(v))$. To this term, agent i^* contributes inefficiency M , agent j' contribute M , and agent j contributes $\left(-\frac{w_{j'}}{w_j}M\right)$, taking the limiting values of $\delta, \gamma, \epsilon \rightarrow 0$. Therefore, the inefficiency term on this profile equals $M + \left(1 - \frac{w_{j'}}{w_j}\right)M > M$, the maximum inefficiency when the weights are equal except i^* (by the arguments just before Lemma 2). Hence, this mechanism cannot achieve the lowest sample inefficiency, which is a contradiction. ■

Lemma 2 and the arguments before it complete the proof of the theorem. ■

Proof of Theorem 3

Proof: Suppose, there exists a deterministic, strategyproof, and neutral mechanism $\langle f, \mathbf{p} \rangle$ that also satisfies $\lim_{n \rightarrow \infty} \sup_{v \in V} T_2^n(\mathbf{p}, v) = 0$. It implies that, at the limit, the mechanism has no budget imbalance, i.e., $\lim_{n \rightarrow \infty} \sup_{v \in V} |\sum_{i=1}^n p_i(v)| = 0$. From the arguments in Footnote 5 (Equations (6) and (7)), we know that f is a neutral affine maximizer and payments are of the form $p_i(v_i, v_{-i}) = h_i(v_{-i}) - \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right)$, $\forall w_i > 0$. We already have the sink mechanism where at least one $w_i = 0$ and the above sum can be made smallest (exactly zero) for every profile $v \in V$. However, that yields a constant upper bound for the term $\lambda \cdot T_1^n(f, v) + (1 - \lambda) \cdot T_2^n(\mathbf{p}, v)$. Hence, we need to consider the case $w_i > 0, \forall i$, which implies that

$$\lim_{n \rightarrow \infty} \sup_{v \in V} \left| \sum_{i=1}^n \left(h_i(v_{-i}) - \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right) \right) \right| = 0.$$

This implies that for every $\epsilon > 0$, there exists $N_\epsilon \in \mathbb{Z}_{\geq 0}$ such that for all $n \geq N_\epsilon$,

$$\left| \sum_{i=1}^n \left(h_i(v_{-i}) - \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right) \right) \right| < \epsilon, \quad \forall v \in V. \quad (17)$$

We show that this identity leads to a contradiction for an appropriately chosen v . Note that this immediately proves the theorem, because if there does not exist any mechanism $\langle f, \mathbf{p} \rangle$ that satisfies the properties mentioned in the theorem statement and makes $\lim_{n \rightarrow \infty} \sup_{v \in V} T_2^n(\mathbf{p}, v) = 0$, then the best possible lower bound is a constant, i.e., $\sup_{v \in V} T_2^n(\mathbf{p}, v) = \Omega(1)$. Therefore, the best lower bound for the spillover factor $\rho_n(f, \mathbf{p})$ is $\Omega\left(\frac{1}{n}\right)$ and this is achievable by the sink mechanism.

We prove this for a set of alternatives $A = \{0, 1\}$ for similar reasons mentioned in Lemma 1. As an illustration of the general proof, let us consider the same argument when $N_\epsilon = 2$. Let the

valuations are $(v_1 + \delta, v_2)$ for alternative 1 and zero otherwise ($\delta > 0$). Also assume that the numbers are such that

$$w_1(v_1 + \delta) + w_2v_2 > 0, \text{ and } w_1v_1 + w_2v_2 < 0.$$

That means, the affine maximizer results in 1 at profile $(v_1 + \delta, v_2)$ and 0 at (v_1, v_2) . The above inequalities can be written concisely as

$$-w_1\delta < w_1v_1 + w_2v_2 < 0. \quad (18)$$

Now by the convergence relation of Equation (17), we have

$$\begin{aligned} \left| h_1(v_2) + h_2(v_1 + \delta) - \frac{w_2}{w_1}v_2 - \frac{w_1}{w_2}(v_1 + \delta) \right| &< \epsilon \\ |h_1(v_2) + h_2(v_1)| &< \epsilon \end{aligned}$$

These inequalities imply¹⁰

$$\begin{aligned} \left| h_2(v_1 + \delta) - \frac{w_2}{w_1}v_2 - \frac{w_1}{w_2}(v_1 + \delta) - h_2(v_1) \right| &< 2\epsilon \\ \Rightarrow \left| \frac{w_2}{w_1}v_2 + \left(h_2(v_1) - h_2(v_1 + \delta) + \frac{w_1}{w_2}(v_1 + \delta) \right) \right| &< 2\epsilon \end{aligned} \quad (19)$$

But this inequality is violated by choosing a large enough δ and large negative v_2 in Equation (18). This is possible to pick since the valuations are picked from $(-M/2, M/2)$ and M is large by definition of ρ_n (Equation (4)). Also note that, the term within parentheses in Equation (19) is independent of v_2 , hence changes in the v_2 will not affect them. Our only assumed relation is Equation (18), and a suitable choice satisfying it violates Equation (19).

The general proof of this theorem extends this idea for any $N_\epsilon = n \geq 2$. Let the agents are numbered in the decreasing order of their weights WLOG, i.e., $w_i \geq w_{i+1}, i = 1, 2, \dots, n-1$. We consider the valuation profile $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1} + \delta, v_n), \delta > 0$ such that

$$-\delta \sum_{i=1}^{n-1} w_i < \sum_{i=1}^n w_i v_i < -\delta \sum_{i=1}^{n-2} w_i \quad (20)$$

The above inequalities imply that the affine maximizer alternative for the profile mentioned above is 1. However, if any agent i 's, $i = 1, 2, \dots, n-1$, valuation changes to v_i from $v_i + \delta$, the alternative changes to 0. We use a generic notation v^S to denote this profile, where S denotes the set of agents such that for all agents $k \in S$, the valuations are v_k , and for all $j \notin S$, the valuations are $v_j + \delta$. Hence, $v^{\{n\}}$ is the profile mentioned before: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1} + \delta, v_n)$ and $v^{\{n-1, n\}}$ is the profile: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1}, v_n)$, for example. Note that, the following term in Equation (17) can be reorganized as

$$\sum_{i=1}^n \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right) = \sum_{i=1}^n \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i(f(v)).$$

¹⁰If $|x + z| < \epsilon$ and $|y + z| < \epsilon$, then $|x - y| = |x + z - (y + z)| \leq |x + z| + |y + z| < 2\epsilon$.

Since, $f(v^{\{n\}}) = 1$, from Equation (17) we have

$$\left| \left(\sum_{i=1}^{n-1} h_i(v_{-i}^{\{n\}}) + h_n(v_{-n}^{\{n\}}) \right) - \left(\sum_{i=1}^n \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) \right| < \epsilon. \quad (21)$$

The idea of the proof is to make a series of substitutions in the first parentheses of the expression above, leaving the terms in the other parentheses unchanged. Note that, the expression in the second parentheses depends on v_n , while the expression $h_n(v_{-n}^{\{n\}})$ does not. The substitutions sequentially eliminate the dependency on v_n from all the terms in the first parentheses, similar to what we did in the two agent case before. This will also increase the RHS of the inequality in Equation (21), but it will be a finite constant factor of ϵ . This leads to a contradiction, since v_n can be chosen arbitrarily large negative by choosing a large positive δ , and still continues to satisfy Equation (20) but violates Equation (21).

The substitutions will involve the term $\sum_{i=1}^{n-1} h_i(v_{-i}^{\{n\}})$ in the first parentheses of Equation (21). Consider the profiles $v^{\{j,n\}}, j = 1, \dots, n-1$. In each of these profiles, $f(v^{\{j,n\}}) = 0$ (due to the choice of $v^{\{n\}}$ in Equation (20)). Hence,

$$\left| \sum_{i=1}^{n-1} h_i(v_{-i}^{\{j,n\}}) + h_n(v_{-n}^{\{j,n\}}) \right| < \epsilon, \quad \forall j \in \{1, \dots, n-1\}. \quad (22)$$

Note that $v_{-i}^{\{i,n\}} = v_{-i}^{\{n\}}$. Hence, we can substitute terms from Equation (22) to the terms in the first parentheses of Equation (21) to get

$$\left| \left(- \sum_{i=1}^{n-1} \sum_{j \neq \{i,n\}} h_j(v_{-j}^{\{i,n\}}) - \sum_{j \neq n} h_n(v_{-n}^{\{j,n\}}) + h_n(v_{-n}^{\{n\}}) \right) - \left(\sum_{i=1}^n \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) \right| < n\epsilon. \quad (23)$$

We now replace the terms $h_j(v_{-j}^{\{i,n\}})$ in the first summation of the first parentheses above. All other terms in that parentheses are h_n functions and, therefore, are independent of v_n . For every $i \neq n$, consider the valuation profiles $v^{\{j,i,n\}}, j \neq i, n$. By Equation (20), $f(v^{\{j,i,n\}}) = 0$, hence, we get an inequality similar to Equation (22):

$$\left| \sum_{k=1}^{n-1} h_k(v_{-k}^{\{j,i,n\}}) + h_n(v_{-n}^{\{j,i,n\}}) \right| < \epsilon, \quad \forall j \neq i, n. \quad (24)$$

Also, note that, $v_{-j}^{\{j,i,n\}} = v_{-j}^{\{i,n\}}$. So, we follow the same procedure to replace the terms $h_j(v_{-j}^{\{i,n\}})$ in Equation (23) to yield a similar inequality where the RHS is replaced by a larger term. Since the number of agents is finite, this process will stop after a finite number of iterations, reducing the terms in the first parentheses only consisting of h_n functions, which are independent of v_n , and the RHS of the inequality being a finite factor $K(n)\epsilon$ (say). This construction shows that the choice of a suitably large δ and negative v_n , which keeps Equation (20) unaffected, can violate the inequality obtained through the iterative procedure described above. Hence the claim is established. ■

Proof of Theorem 4

Proof: Assume $m = n + 1$. The proof generalizes to any $m > n$. Consider the valuation profile $v_i = (v_i(a_1), \dots, v_i(a_n), v_i(a_{n+1}))$ where $v_i(a_i) = -M/2 + \epsilon/2$, $v_i(a_{n+1}) = M/2 - \epsilon$ and $v_i(a_j) = M/2 - \epsilon/2$, $\forall j \neq i, n+1$, and $\epsilon > 0$ is arbitrarily small, $i \in N$. This profile is possible to construct since $m > n$. Clearly, the efficient alternative is a_{n+1} , but if any agent i is picked as a sink, the alternative changes to a_i , which has inefficiency of $M - \epsilon$. Therefore the expected sample inefficiency for any generalized sink mechanism is $\frac{1}{nM}(M - \epsilon)$. Taking $\epsilon \rightarrow 0$ proves the theorem. ■

Proof of Theorem 5

Proof: Suppose, for m_2 alternatives the valuation profile v^* yields the worst inefficiency $r_{n,m_2}^M(f)$. Clearly, we can append the other alternatives when we increase the number of alternatives to m_1 with values arbitrarily close to $-M/2$ so that they never change the optimal alternative. Hence the inefficiency cannot decrease. ■

Proof of Theorem 6

Proof: Consider an arbitrary agent i . If agent i is chosen as a sink, the maximum *absolute* inefficiency that it can produce is M (by the same argument as in Equation (16), and we refer to the unnormalized difference term $\sum_{i \in N} v_i(a^*(v)) - \sum_{i \in N} v_i(f(v))$ as the absolute inefficiency). Say the efficient alternative is a . This inefficiency is achieved when the sum of valuations of the agents other than i at the other alternative b is just higher than that of those agents at a , i.e., $\sum_{j \neq i} v_j(b) = \sum_{j \neq i} v_j(a) + \epsilon$, where $\epsilon > 0$ and small, and also the difference in valuations of agent i at these two alternatives is maximum, i.e., $v_i(a) - v_i(b) = M - \delta$, where $\delta > 0$ and small. This implies that, without i , the population is almost equally divided among the alternatives a and b with a marginal bias to b and agent i is ‘maximally’ in favor of a . The difference $v_i(a) - v_i(b) = M - \delta$ is achieved only when the values are close to $v_i(a) = \frac{M}{2} - \frac{\delta}{2}$ and $v_i(b) = -\frac{M}{2} + \frac{\delta}{2}$, since all valuations must lie within $(\frac{M}{2}, -\frac{M}{2})$. Agent i , therefore, is *pivotal* for making the decision in favor of a .

Now, we claim that there cannot be more than $\lceil \frac{n}{2} \rceil$ such pivotal agents i . Suppose for contradiction that there are $> \lceil \frac{n}{2} \rceil$ such pivotal agents. We present the argument for $\lceil \frac{n}{2} \rceil + 1$ for brevity, but it generalizes to the case with more number of pivotal agents. For each of these agents k , $v_k(a)$ is arbitrarily close to $\frac{M}{2}$ and $v_k(b)$ is arbitrarily close to $-\frac{M}{2}$. Therefore, if any of them, say agent k^* , is chosen as a sink, there are $\lceil \frac{n}{2} \rceil$ other similar agents who still make the decision of the mechanism in favor of a (since the sum valuation for a will be larger than b by an unsurmountable value $\approx \lceil \frac{n}{2} \rceil M$) irrespective of the valuation profiles of the other agents. This implies that k^* is not a pivotal agent, which is a contradiction.

The mechanism chooses each sink agent with probability $\frac{1}{n}$. Therefore, the expected inefficiency can at most be $\frac{1}{n} \cdot \lceil \frac{n}{2} \rceil \cdot M$. Divide this by nM to get the expected sample inefficiency. It is easy to see that this bound is tight. The valuation profile that achieves this bound has $\lceil \frac{n}{2} \rceil$ agents having valuations $v(a) \approx M/2$, $v(b) \approx -M/2$ and the rest of the agents have the reverse valuations. ■

Proof of Theorem 8

Proof: For $n = m = 2$, either the two agents prefer the same candidate or opposing candidates. (Ties are broken in favor of a_1 , say.) If they agree, any probability of the generalized sink mech-

anism yields zero absolute inefficiency, since the efficient alternative will be chosen irrespective of which agent is the sink. Note that the payments are always zero for two-agent generalized sink mechanisms. When the agents oppose, we claim that in every opposing profile, a strategyproof generalized sink mechanism must have the same probability of picking the sinks. Suppose not, that is, for a specific generalized sink mechanism $g : V \rightarrow \Delta N$, the probabilities of picking the sinks are different in profiles $v = (v_1, v_2)$ and $v' = (v'_1, v'_2)$, i.e., $g(v) \neq g(v')$. Consider the transition: $v = (v_1, v_2) \rightarrow (v'_1, v_2) \rightarrow (v'_1, v'_2) = v'$. The sink-picking probabilities g must have changed in at least one of these two transitions, that is, either $g(v_1, v_2) \neq g(v'_1, v_2)$ or $g(v'_1, v_2) \neq g(v'_1, v'_2)$. But this is a contradiction to strategyproofness since at least one agent will misreport in that profile pair. She will prefer to increase the probability of the other agent becoming sink so that her favorite candidate has higher probability of being selected, which increases her utility since payment is zero. For example, suppose in the first transition, the probability of agent 1 being sink is higher in the profile (v_1, v_2) , and consequently, probability of agent 2 being sink is lower. Then agent 1 will misreport her valuation to v'_1 . Now, among all fixed probability distributions, $(0.5, 0.5)$ gives the minimum absolute inefficiency which is $\frac{1}{2}$. ■

Proof of Theorem 9

The problem of Equation (5) can be written in a standard LP representation as follows. The notation $f_a(v)$ denotes the probability of picking alternative a by the randomized mechanism f when the valuation profile is v .

$$\min_{f, \mathbf{p}} \quad \ell \tag{25}$$

s.t.

$$\begin{aligned} & [v_1(a) \cdot f_a(v_1, v_2) + v_1(b) \cdot f_b(v_1, v_2) - p_1(v_1, v_2)] - [v_1(a) \cdot f_a(v'_1, v_2) \\ & \quad + v_1(b) \cdot f_b(v'_1, v_2) - p_1(v'_1, v_2)] \geq 0, \quad \forall v_1, v'_1, v_2, \quad \textbf{Agent 1, SP} \\ & [v_2(a) \cdot f_a(v_1, v_2) + v_2(b) \cdot f_b(v_1, v_2) - p_2(v_1, v_2)] - [v_2(a) \cdot f_a(v_1, v'_2) \\ & \quad + v_2(b) \cdot f_b(v_1, v'_2) - p_2(v_1, v'_2)] \geq 0, \quad \forall v_1, v_2, v'_2, \quad \textbf{Agent 2, SP} \end{aligned}$$

$$\begin{aligned} & f_a(v) + f_b(v) = 1, \quad \forall v \in V, \quad \textbf{SCF} \\ & p_1(v) + p_2(v) = 0, \quad \forall v \in V, \quad \textbf{Budget Balance} \end{aligned}$$

$$\begin{aligned} & \ell + (v_1(a) + v_2(a)) \cdot f_a(v) + (v_1(b) + v_2(b)) \cdot f_b(v) \\ & \geq \max_{x \in \{a, b\}} (v_1(x) + v_2(x)), \quad \forall v \in V, \quad \textbf{Max Inefficiency} \\ & f_a(v), f_b(v) \geq 0, \quad \forall v \in V \end{aligned}$$

Proof: For $k = 3$, each agent has $3^2 = 9$ valuations, since the number of alternatives is 2, and therefore, the number of valuation profiles is 81. The optimization variables are

$$\mathbf{x} := (f_a(v^0), f_b(v^0), \dots, f_a(v^{80}), f_b(v^{80}), p_1(v^0), p_2(v^0), \dots, p_1(v^{80}), p_2(v^{80}), \ell)^\top.$$

Here the 81 valuation profiles are indexed from 0 to 80 and are denoted by the superscripts. Hence there are $81 \times 4 + 1 = 325$ variables to the discretized relaxation of the primal problem of Equation (25). However, we can significantly reduce the number of variables using the symmetry of the LP. The symmetry that we consider are *anonymity*, i.e., the SCF alternative is invariant to the permutation of the agents, and the payments are permuted according to the permutation of the agents (defined below), and *neutrality*, i.e., the relabeling of the alternatives changes the alternative according to the same relabeling (Definition 4).

DEFINITION 9 (Anonymity) *A mechanism $\langle f, \mathbf{p} \rangle$ is anonymous if for every permutation of the agents λ we have*

$$f(\lambda(v)) = f(v) \quad \text{and} \quad p_{\lambda(i)}(\lambda(v)) = p_i(v), \quad \forall v \in V, \forall i \in N.$$

Note that, similar to Definition 4, we have overloaded the notation of λ for the valuation profile to denote that the profile where the agents are permuted according to λ .

LEMMA 3 *For every strategyproof, budget-balanced, randomized mechanism that achieves the minimum absolute inefficiency, there exists an anonymous, neutral, strategyproof, budget-balanced, randomized mechanism that achieves the same absolute inefficiency.*

Proof: We prove this for two agents and two alternatives. The same argument generalizes to any number of agents and alternatives. Consider an optimal solution of the optimization problem of Equation (25). This yields a solution \mathbf{x}^* (say). Suppose, we relabel the agents 1 and 2 by swapping their identities, which changes the valuation profiles accordingly. For example, now the payment $p_1(v_1, v_2)$ is swapped with $p_2(v_2, v_1)$. We keep the SCF alternatives identical, i.e., $f_a(v_1, v_2) = f_a(v_2, v_1)$. Now consider the resulting vector of variables $\mathbf{x}_{\text{AGENT-SWAP}}^*$. Note that, this permutation of the variables reorders the set of constraints in Equation (25). The SP constraints of agent 1 now becomes the SP constraints of agent 2 and vice-versa. SCF constraints remain identical, budget balance constraints are reordered but same, and the max-inefficiency constraints are also reordered. Hence $\mathbf{x}_{\text{AGENT-SWAP}}^*$ is a feasible solution of the LP (Equation (25)) and since $\mathbf{x}_{\text{AGENT-SWAP}}^*$ and \mathbf{x}^* has the same value for ℓ , $\mathbf{x}_{\text{AGENT-SWAP}}^*$ is an optimal solution of the LP (Equation (25)).

Similarly, we swap the alternatives a and b and the valuations accordingly to obtain a different reordered vector $\mathbf{x}_{\text{ALT-SWAP}}^*$. This relabeling of the alternatives again reorders all the constraints in a different way than the earlier case, with ℓ remaining same in both these cases. In a similar way as before, we argue that $\mathbf{x}_{\text{ALT-SWAP}}^*$ is an optimal solution of the LP (Equation (25)).

Now, we swap both the alternatives and agents to obtain $\mathbf{x}_{\text{AGENT-ALT-SWAP}}^*$ which reorders the constraints in a two-fold manner, but the last variable of this vector remains ℓ as before, and therefore it is also an optimal solution of the LP (Equation (25)).

Now, we have 4 optimal solutions given the original optimal solution \mathbf{x}^* , which are complementary to each other in terms of agents and alternatives, but all of them are strategyproof, budget-balanced, randomized mechanisms (since they are feasible solutions of Equation (25)). Consider the average of all these solutions:

$$\mathbf{x}^{A,N} = \frac{1}{4}(\mathbf{x}^* + \mathbf{x}_{\text{AGENT-SWAP}}^* + \mathbf{x}_{\text{ALT-SWAP}}^* + \mathbf{x}_{\text{AGENT-ALT-SWAP}}^*).$$

By construction, $\mathbf{x}^{A,N}$ is anonymous and neutral, but this is also another optimal solution of the LP (Equation (25)) (since the set of constraints is convex). Hence, we have proved the lemma for two agents and two alternatives. For n agents and m alternatives, we consider all $n!$ and $m!$ possible permutations of the agents and alternatives respectively and take the mean of them to obtain our resulting optimal solution that is both anonymous and neutral. ■

Hence, it is WLOG to consider neutral and anonymous mechanisms to solve the optimization problem of Equation (25). This reduces the number of variables in the primal problem, since for valuations that are either agent permuted or alternative permuted or both permuted version of a valuation profile we have already considered, we can replace their constraints with the already considered variables. We write the coefficient matrix of the constraint set of the LP in Equation (25) denoted by A as follows.

$$\begin{array}{c}
 \begin{array}{cccccccccccccccc}
 f_a(v^0) & f_b(v^0) & \dots & \dots & f_a(v^{80}) & f_b(v^{80}) & p_1(v^0) & p_2(v^0) & \dots & \dots & p_1(v^{80}) & p_2(v^{80}) & \ell
 \end{array} \\
 \left(\begin{array}{cccccccccccccccc}
 v_1^0(a) & v_1^0(b) & -v_1^8(a) & -v_1^8(b) & 0 & 0 & -1 & 0 & 1 & 0 & 0 & 0 & 0 \\
 v_2^0(a) & v_2^0(b) & -v_2^1(a) & -v_2^1(b) & 0 & 0 & 0 & -1 & 0 & 1 & 0 & 0 & 0 \\
 0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & \dots & \dots & 0 & 0 & 0 \\
 0 & 0 & -v_2^{79}(a) & -v_2^{79}(b) & v_2^{80}(a) & v_2^{80}(b) & 0 & 0 & 0 & 1 & 0 & -1 & 0 \\
 0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & \dots & \dots & 0 & 0 & 0 \\
 \hline
 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 \hline
 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & \dots & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\
 \hline
 w^0(a) & w^0(b) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
 0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
 0 & 0 & 0 & 0 & w^{80}(a) & w^{80}(b) & 0 & 0 & 0 & 0 & 0 & 0 & 1
 \end{array} \right) \quad (26)
 \end{array}$$

Where $w^p(x) = v_1^p(x) + v_2^p(x)$, $x \in \{a, b\}$, and p denotes the profile index. The header of the matrix shows the primal variables. The sections showed in dotted lines corresponds to the strategyproofness, valid SCF, budget balance, and maximum inefficiency constraints respectively. The RHS of the constrained inequalities of the LP is a vector \mathbf{b} that looks as follows:

$$\mathbf{b} := (0, \dots, 0, \mathbf{1}_{|V|}, \mathbf{0}_{|V|}, \max_x w^0(x), \dots, \max_x w^{80}(x))^\top.$$

Denoting the cost vector of the LP as, $\mathbf{c} := (\mathbf{0}_{4|V|}, 1)^\top$, we can represent the LP of Equation (25) in the standard form:

$$\begin{array}{ll}
 \text{primal} & \min_{\mathbf{x}} \quad \mathbf{c}^\top \mathbf{x} \\
 & \text{s.t.} \quad A\mathbf{x} \geq \mathbf{b} \\
 \text{dual} & \max_{\mathbf{y}} \quad \mathbf{b}^\top \mathbf{y} \\
 & \text{s.t.} \quad \mathbf{y}^\top A \leq \mathbf{c}^\top
 \end{array} \quad (27)$$

Our goal is to provide a lower bound of the optimal value of the primal. Hence, we consider its dual, and provide a feasible solution. By weak duality lemma, the value of the dual objective at

that feasible point will be a lower bound of the primal. The dual variables represented by \mathbf{y} consists of $(\lambda, \gamma, \mu, \delta)$. The λ variables refer to the dual variables corresponding to the strategyproofness constraints, and we denote the dual variable that represent the strategyproofness of agent i between the profiles v^k and v^l by λ_{i,v^k,v^l} . By this representation, we consider only such pairs of profiles v^k and v^l where only agent i 's valuation changes. The γ variables are the dual variables corresponding to the constraint that the SCF must add to unity, and we denote the dual variable corresponding to value profile v as γ_v . Dual variables μ and δ corresponds to the budget balance and the maximum inefficiency constraints. Since SCF and budget balance constraints are equalities, γ and μ are unrestricted, while λ and δ are non-negative. Additionally, in the primal problem the payment variable p_i 's were unrestricted, hence in the dual the corresponding constraints are equalities.

We now provide a dual feasible solution, which is represented with respect to the reduced set of dual variables. Using symmetry according to Lemma 3, we reduce the number of valuation profiles. We number the profiles from 0 to 80 in the following way: for the valuation $(-0.5, -0.5)$ of agent 1, all possible valuations of agent 2 from $(-0.5, -0.5)$ to $(0.5, 0.5)$ (9 profiles) are listed, and then the valuation of agent 1 is moved to $(-0.5, 0)$. Due to symmetry, setting a primal variable $f_a(v)$ to a certain value also fixes 3 other variables that are agent-swapped or alternative-swapped or both-swapped versions of this variable. Denote the reduced set of valuation profiles by V_R . This also reduces the dual variables γ, μ, δ from 81 to 27 independent variables. However, for the λ variables we need to list all of them since they correspond to constraints that involve two valuation profiles. Consider the following set of dual variables (numbers of v and v' correspond to the valuation profile numbers in the listing discussed above):

i	v	v'	λ				
1	11	2	4/7	v	γ	v	μ
1	12	21	4/7				
1	30	66	1/14	2	2/7	6	-1/14
1	52	16	2/7	6	3/28	7	1/7
1	57	30	1/7	7	1/7	7	1/7
1	60	33	3/14	8	2/7	8	4/7
1	68	32	2/7	10	-1/28	11	-4/7
1	78	15	4/7	11	2/7	12	-4/7
2	12	16	4/7	14	-3/14	14	3/14
2	14	10	1/14	19	-4/7	24	2/7
2	18	26	4/7	24	-1/7		
2	20	19	3/14				
						v	δ
						2	2/7
						6	1/7
						11	4/7

All other entries of the variables that are not shown in the list above are zero. Note that for only the λ variables, the valuation profiles listed go beyond the index 27, but for all other dual variables they are represented by the 27 independent variables listed in V_R .

We claim that this is a feasible solution of the dual. The proof requires an exhaustive verification for each of the inequalities in the constraint set of the dual. However, we provide a few cases to give an insight how this example is picked. Consider, the variables $\lambda(1, 52, 16) = 2/7$ and $\lambda(1, 68, 32) = 2/7$. Note that, $v^{24} = ((0, 0.5), (0.5, 0)) = v^{52}$ and $v^{68} = ((0.5, 0), (0, 0.5))$ is an alternative swapped version of v^{24} . Also, none of the other variables involve any agent or alternative or both swap of this profile in the example we gave. Therefore, now we need to concentrate on the column $f_b(v^{24})$ in the matrix of Equation (26). Note that the matrix of Equation (26) is also

reduced on the column and the rows. On the column, each of the f and p columns are reduced to $|V_R|$, and on the rows, only the strategyproofness constraints retain the original number, but the SCF, budget balance and maximum inefficiency constraints reduce to $|V_R|$ in size. Carrying out the product with the terms we get $0.5 \times 2/7 + 0 \times 2/7 = 1/7$. While inspecting other variables, we find $\gamma(v^{24}) = -1/7$. Hence the sum of the products on the column $f_b(v^{24})$ gives $1/7 - 1/7 = 0$ which satisfies the inequality. This is not an isolated case, in all the columns $f_x(v)$ (the numbers of such variables are reduced because of symmetry), the examples are chosen such that the sum of the non-zero products in the SP constraints section and one non-zero product in the SCF constraints section add up to a non-positive number (for example, repeat the same argument for v^{78} and v^{19}).

Similarly, consider the column $p_2(v^{12})$: the variable $\lambda(2, 12, 16) = 4/7$ gets multiplied with -1 in this column since the constraint for agent 2 in the profile v^{12} gives a -1 coefficient for $p_2(v^{12})$. However, the variable $\mu(v^{12}) = -4/7$ which is multiplied with 1 in this column, and we can inspect that no other product is non-zero on this column. Hence the sum of the products is $-8/7$ which is non-positive, and satisfies the dual constraint.

The easiest thing to verify is the last column, where the sum of the δ_v for the reduced set of v 's add to unity ($2/7 + 1/7 + 4/7$). Therefore, the example provided is a dual feasible solution. We compute the objective value of the solution:

$$\begin{aligned} & \sum_{v \in V_R} \gamma_v + \sum_{v \in V_R} \delta_v \max_{x \in \{a, b\}} w^v(x) \\ &= \frac{2}{7} + \frac{3}{28} + \frac{1}{7} + \frac{2}{7} - \frac{1}{28} + \frac{2}{7} - \frac{3}{14} - \frac{4}{7} - \frac{1}{7} + 0.5 \times \frac{4}{7} \\ &= \frac{1}{7} \end{aligned}$$

This completes the proof of the theorem. ■